

Fast integral equation methods for the modified Helmholtz equation

Mary-Catherine A. Kropinski^{a,*}, Bryan Quaife^{a,*}

^a*Department of Mathematics, Simon Fraser University, Burnaby, British Columbia,
Canada V5A 1S6*

Abstract

We present a collection of integral equation methods for the solution to the two-dimensional, modified Helmholtz equation, $u(\mathbf{x}) - \alpha^2 \Delta u(\mathbf{x}) = 0$, in bounded or unbounded multiply-connected domains. We consider both Dirichlet and Neumann problems. We derive well-conditioned Fredholm integral equations of the second kind, which are discretized using high-order, hybrid Gauss-trapezoid rules. Our fast multipole-based iterative solution procedure requires only $O(N)$ or $O(N \log N)$ operations, where N is the number of nodes in the discretization of the boundary. We demonstrate the performance of the methods on several numerical examples.

Keywords: Fast multipole method, Gaussian quadrature, Modified Helmholtz equation, integral equations, Yukawa potential.

1. Introduction

A variety of important problems in science and engineering involve the solution to the problem

$$u(\mathbf{x}) - \alpha^2 \Delta u(\mathbf{x}) = f(\mathbf{x}), \quad (1)$$

subject to appropriate boundary conditions. This equation is called the modified Helmholtz equation. It appears, for example, in the semi-implicit temporal discretization of the heat or the Navier-Stokes equations [7] (here

*Supported in part by Natural Sciences and Engineering Research Council of Canada Grant RGPIN 203326

Email address: mkropins @ cs.sfu.ca (Mary-Catherine A. Kropinski)

α^2 would be proportional to the time step), and in the linearized Poisson-Boltzmann equation. The underlying free-space Green's function is referred to as the Yukawa or screened-Coulomb potential.

In [4], Cheng et al. present a fast direct solver for (1) in two dimensions on the unit square. The solution is expressed as a volume potential, and the direct solver is accelerated using a new version of the fast multipole method [6, 10]. The solver is fully adaptive and the computational costs are comparable to those of FFT-based methods. Our aim, here, is to complement this work. In order to solve (1) in more general domains, solutions to

$$u(\mathbf{x}) - \alpha^2 \Delta u(\mathbf{x}) = 0, \quad (2)$$

with prescribed boundary conditions are required. Here, we present integral equation methods to solve (2) subject to Neumann or Dirichlet boundary conditions in multiply-connected domains, which may be bounded or unbounded in extent.

Representing (2) as an integral equation is a natural choice, and computational methods based on an integral equation formulation can have significant advantages over conventional finite difference or finite element techniques. Integral equation methods are naturally adaptive, easily allow for high-order approximations, and can handle arbitrarily complex boundaries. However, the discretization of most integral operators yield dense matrices, and in the absence of fast algorithms to solve these systems, integral equation methods are rarely competitive.

We present fast-multipole accelerated methods for solving integral equation representations of (2). This work is meant to join a growing collection of fast-multipole-accelerated integral equation methods for linear, elliptic operators [4, 5, 9, 8]. The fast multipole method (FMM) was first introduced by Greengard and Rokhlin as an efficient way to evaluate the Coulomb potential due to a collection of charged particles [3]. It has since been extended to handle different potentials, in both two and three dimensions. Of particular interest here are the FMM methods that have been developed to compute volume potentials for the Yukawa potential in two dimensions [4], and particle interactions for the three-dimensional Yukawa potential [6]. Our methods have the following key elements:

- Well-conditioned Fredholm integral equations of the second kind are formulated.

- The integral equations are discretized with tailored, high-order accurate hybrid Gauss-trapezoid rules.
- The resulting linear systems is solved using a GMRES scheme [18].
- The FMM for the two-dimensional Yukawa potential is exploited to compute the matrix-vector products in the iterative solution procedure.

With N points in the discretization of the boundary, our methods require only $O(N)$ or $O(N \log N)$ operations.

There has been relatively little work done on developing integral equation methods for the modified Helmholtz equation. The majority of the work appears to be focused on the linearized Poisson-Boltzmann equation, of which [13, 14, 15] are recent examples. Here, the interest is typically to perform an electrostatic analysis over a single complicated molecular surface, with the electrostatic potential satisfying certain jump conditions across the boundary. Our intended application is to couple our work with that presented in [4] in order to solve the equations that arise from the temporal discretization of the heat equation, for example. Thus, we are more concerned with developing a general-purpose solver.

There is a growing body of literature on fast algorithms for boundary integral equation methods. Nishimura presents a thorough review in [17]. Recent technological advances include using a “kernel-independent” fast-summation algorithm, as discussed by Ying et al. in [19]. Their method results in an $O(N^{3/2})$ scheme. Another exciting new development includes fast direct solvers [12, 16], which directly construct a compressed factorization of the inverse of the matrix arising from the discretization of certain boundary integral equations. These methods are particularly attractive to problems in which it is of interest to solve the equations with multiple right-hand sides.

While significant progress has already been made in developing efficient integral equation methods for linear elliptic equations, similar treatment of time-dependent equations has lagged behind. Also, the application of these techniques to practical problems of interest in the broader scientific community has been limited. The reason seems clear; the majority of partial differential equations that arise from modelling real-world phenomena are both time dependent and nonlinear. In [7], Greengard and Kropinski argued that after both discretizing in time and applying a suitable linearization, the Navier-Stokes equations can be rendered amenable to treatment by integral equation methods. However, in order to realize this, certain “building-block”




Figure 1: A bounded multiply-connected domain D . The outer boundary is denoted by Γ_0 (which is absent in unbounded domains) and the interior component curves by $\Gamma_1, \dots, \Gamma_M$. The unit normal \mathbf{n} points out of D on each component curve.

tools are still needed. Our work is meant to provide one such tool and bring us one step closer to making integral equation methods an attractive alternative to conventional finite element and finite difference methods for many large-scale problems in science and engineering.

The paper is organized as follows: In section 2, we discuss the potential theory for the modified Helmholtz operator and the corresponding integral equation formulations. In section 3, we present our numerical methods, with the fast-multipole method being briefly outlined in section 4. Numerical examples are presented in Section 5.

2. Potential theory and integral equation formulations

To fix notation, let us consider a domain D with boundary Γ which is M or $(M + 1)$ -ply connected (Figure 1). We are interested in solving (2), subject to both Dirichlet and Neumann boundary conditions. In both cases, we start with the free-space Green's function $G(\mathbf{y} - \mathbf{x}) = G(y_1 - x_1, y_2 - x_2)$ for the operator $1 - \alpha^2 \Delta$,

$$G(\mathbf{x}) = \frac{1}{2\pi\alpha^2} K_0 \left(\frac{|\mathbf{x}|}{\alpha} \right),$$

where K_0 is the zeroth-order modified Bessel's function of the second kind. The Dirichlet problem is considered in section 2.1 and the Neumann problem in 2.2.

2.1. The Dirichlet Problem

Consider the following Dirichlet problem:

$$\begin{aligned} u(\mathbf{x}) - \alpha^2 \Delta u(\mathbf{x}) &= 0, & \mathbf{x} \in D, \\ u(\mathbf{x}) &= f(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{aligned} \quad (3)$$

If D is unbounded, the appropriate decay conditions for the solution is $u(\mathbf{x}) \rightarrow 0$ as $|\mathbf{x}| \rightarrow \infty$. We seek the solution $u(\mathbf{x})$ in the form of a double layer potential,

$$u(\mathbf{x}) = \frac{1}{2\pi\alpha^2} \int_{\Gamma} \frac{\partial}{\partial n_{\mathbf{y}}} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \sigma(\mathbf{y}) ds_{\mathbf{y}}, \quad \mathbf{x} \in D, \quad (4)$$

where $\sigma(\mathbf{y})$ is the value of the unknown density at the boundary point \mathbf{y} , and $\partial/\partial n_{\mathbf{y}}$ represents the outward normal derivative at the point \mathbf{y} .

In order to derive the form of the integral equation based on the representation (4), we must understand the behaviour of the potential on the boundary Γ . We note that as $z \rightarrow 0$,

$$K_0(z) \sim -\log(z) + Q(z),$$

where $Q(z)$ is a polynomial in z [1]. Thus, (4) behaves asymptotically as if it has a logarithmic kernel, and we note for future reference that

$$\lim_{\substack{\mathbf{y} \rightarrow \mathbf{x} \\ \mathbf{x}, \mathbf{y} \in \Gamma}} \frac{\partial}{\partial n_{\mathbf{y}}} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) = -\frac{1}{2} \kappa(\mathbf{x}), \quad (5)$$

where $\kappa(\mathbf{x})$ denotes the curvature of Γ at the point \mathbf{x} .

The jump relations for potentials of the logarithmic kernel are well known, and therefore, for any point \mathbf{x} on the boundary Γ ,

$$\begin{aligned} \lim_{\substack{\mathbf{x}' \rightarrow \mathbf{x} \\ \mathbf{x}' \in D}} u(\mathbf{x}') &= -\frac{1}{2\alpha^2} \sigma(\mathbf{x}) + \frac{1}{2\pi\alpha^2} \int_{\Gamma} \frac{\partial}{\partial n_{\mathbf{y}}} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \sigma(\mathbf{y}) ds_{\mathbf{y}}, \\ \lim_{\substack{\mathbf{x}' \rightarrow \mathbf{x} \\ \mathbf{x}' \in D^c}} u(\mathbf{x}') &= \frac{1}{2\alpha^2} \sigma(\mathbf{x}) + \frac{1}{2\pi\alpha^2} \int_{\Gamma} \frac{\partial}{\partial n_{\mathbf{y}}} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \sigma(\mathbf{y}) ds_{\mathbf{y}}. \end{aligned}$$

Using the first limit in the preceding expression, and matching to the boundary condition (3), we obtain an integral equation for the layer density σ :

$$\sigma(\mathbf{x}) - \frac{1}{\pi} \int_{\Gamma} \frac{\partial}{\partial n_{\mathbf{y}}} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \sigma(\mathbf{y}) ds_{\mathbf{y}} = -2\alpha^2 f(\mathbf{x}). \quad (6)$$

2.2. The Neumann Problem

Consider the following Neumann problem:

$$\begin{aligned} u(\mathbf{x}) - \alpha^2 \Delta u(\mathbf{x}) &= 0, & \mathbf{x} \in D, \\ \frac{\partial u(\mathbf{x})}{\partial n} &= g(\mathbf{x}), & \mathbf{x} \in \Gamma. \end{aligned}$$

We seek the solution $u(\mathbf{x})$ in the form of a single layer potential

$$u(\mathbf{x}) = \frac{1}{2\pi\alpha^2} \int_{\Gamma} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \sigma(\mathbf{y}) ds_{\mathbf{y}}.$$

In order to satisfy the boundary conditions, the density $\sigma(\mathbf{x})$ must satisfy the integral equation

$$\sigma(\mathbf{x}) + \frac{1}{\pi} \int_{\Gamma} \frac{\partial}{\partial n_{\mathbf{x}}} K_0 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \sigma(\mathbf{y}) ds_{\mathbf{y}} = 2\alpha^2 g(\mathbf{x}). \quad (7)$$

We note that from (5), the kernels in the integral equations (6) and (7) are bounded and continuous along Γ . In addition, there are no nontrivial homogeneous solutions. Therefore, by the Fredholm alternative, (6) and (7) have unique solutions for any integrable data $f(\mathbf{x})$ or $g(\mathbf{x})$.

In summary, equations (6) and (7) can be written in the general form:

$$\sigma(\mathbf{x}) + \frac{1}{\pi} \int_{\Gamma} K(\mathbf{y}, \mathbf{x}) \sigma(\mathbf{y}) ds_{\mathbf{y}} = F(\mathbf{x}). \quad (8)$$

In the case of the Dirichlet problem, we have

$$\begin{aligned} K(\mathbf{y}, \mathbf{x}) &= \frac{1}{\alpha} K_1 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \frac{\mathbf{y} - \mathbf{x}}{|\mathbf{y} - \mathbf{x}|} \cdot \mathbf{n}_{\mathbf{y}}, \\ F(\mathbf{x}) &= -2\alpha^2 f(\mathbf{x}). \end{aligned}$$

In the case of the Neumann problem, we have

$$\begin{aligned} K(\mathbf{y}, \mathbf{x}) &= \frac{1}{\alpha} K_1 \left(\frac{|\mathbf{y} - \mathbf{x}|}{\alpha} \right) \frac{\mathbf{y} - \mathbf{x}}{|\mathbf{y} - \mathbf{x}|} \cdot \mathbf{n}_{\mathbf{x}}, \\ F(\mathbf{x}) &= 2\alpha^2 g(\mathbf{x}). \end{aligned}$$

3. Numerical methods

We assume each component curve Γ_k , $k = 1, \dots, M$, is parametrized by $\mathbf{y}^k(\alpha)$, where $\alpha \in [0, 2\pi)$. Similarly, $\sigma^k(\alpha)$ refers to the restriction of the density σ on Γ_k . On each contour Γ_k , we are given N points equispaced with respect to α . Thus the mesh spacing is $h = 2\pi/N$, and the total number of discretization points is NM . Associated with each such point, denoted by \mathbf{y}_j^k , is an unknown density σ_j^k . The derivative $\partial\mathbf{y}/\partial\alpha$ is denoted by \mathbf{w}^k , and the derivative values \mathbf{w}_j^k are obtained pseudospectrally.

The presence of the logarithmic singularity in the integral operator causes the spectrum of kernels based on the Yukawa potential to decay slowly. Applying the straightforward Nyström discretization based on the trapezoidal rule would result in a significant loss of accuracy. Instead, we use hybrid gauss-trapezoidal quadrature rules developed by Alpert [2] which are tailored for integrands with logarithmic singularities. These quadratures are of order $h^p \log h$. The order p determines a set of nodes u_n and weights v_n , $n = 1 \dots l$. These nodes and weights are used for the quadrature within the interval $\alpha \in [\alpha_j - ha, \alpha_j + ha]$, on Γ_k (a and l are also determined by p). Outside of this interval, the trapezoid rule is used. Applying this quadrature to (8) yields

$$\begin{aligned} \sigma_j^k + \frac{h}{\pi} \left\{ \sum_{\substack{m=0 \\ m \neq k}}^M \sum_{n=1}^N K(\mathbf{y}_n^m, \mathbf{y}_j^k) |\mathbf{w}_n^m| \sigma_n^m + \sum_{n=j+a}^{N+j-a} K(\mathbf{y}_n^k, \mathbf{y}_j^k) |\mathbf{w}_n^k| \sigma_n^k \right\} \\ + \frac{1}{\pi} \sum_{\substack{n=-l \\ n \neq 0}}^l u_{|n|} K(\mathbf{y}_{j+\frac{n}{|n|}v_{|n|}}^k, \mathbf{y}_j^k) |\mathbf{w}_{j+\frac{n}{|n|}v_{|n|}}^k| \sigma_{j+\frac{n}{|n|}v_{|n|}}^k = F_j^k. \end{aligned} \quad (9)$$

In the second sum, we invoke periodicity of all functions on Γ_k , or equivalently, $j + N_k = j$. In the final sum, we are required to know values of σ intermediate to the nodal values at $\alpha = \alpha_j \pm hv_{|n|}$. In these cases, we use Fourier interpolation.

Equation (9) is a dense $MN \times MN$ linear system that must be solved for the unknowns σ_j^k . In our implementation, (9) is solved iteratively, using GMRES [18]. The bulk of the work at each iteration lies in applying the linear system to a vector. Done directly, this would require $O(N^2M^2)$ operations for each iteration. This cost can be reduced to $O(NM)$ by using the adaptive fast multipole method, the details of which are outlined in the next section.

Since the number of iterations needed to solve a Fredholm equation of the second kind to a fixed precision is bounded independent of the system size N (holding the geometry fixed), the total number of operations required to solve (9) is $O(NM)$.

4. The fast multipole method

Consider a collection of N particles or “sources” in \mathbb{R}^2 , $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$, together with corresponding source strengths q_1, q_2, \dots, q_N . We are interested in efficiently evaluating the following:

$$\Phi(\mathbf{x}_j) = \sum_{\substack{i=1 \\ i \neq j}}^N q_i K_0 \left(\frac{|\mathbf{x}_j - \mathbf{x}_i|}{\alpha} \right), \quad j = 1, \dots, N. \quad (10)$$

The fast multipole method was developed to evaluate fields based on the Coulomb potential in two dimensions [3, 11]. The FMM has since been extended to include other potential functions, such as the three-dimensional Yukawa potential [6]. An FMM developed to evaluate volume potentials based on the two-dimensional Yukawa potential appears in [4]; we use this work as the basis for our “particle-to-particle” FMM required to evaluate (9). Our FMM is identical in structure to the one presented in [4], with only minor modification. For more detail, we refer the reader to this work and others [6, 10] which discusses the use of exponential expansions. Here, we present a minimal sketch of the FMM algorithm as it applies to our problem at hand.

The FMM uses an adaptive quad-tree structure in order to superimpose a hierarchy of refinements on the computational domain. We imbed the geometry inside a unit square S , which is considered to be grid level 0. Grid level $l + 1$ is obtained recursively by subdividing each square s at level l into four equal parts, which are the “children” of s . Adaptivity is achieved by not requiring the same number of levels of subdivision in all regions of S . The basic idea of the FMM is that for each particle, contributions from “nearby” (neighbour) particles to the potential field are handled directly via (10), while “far-field” (non-neighbour) interactions are handled using multipole and/or related expansions.

The first step in the FMM is to form the multipole expansions for all of the nodes (boxes) in the quad-tree structure. The following theorem follows from Graf’s addition theorem, and corresponds to Theorem 3 in [4] with modifications made for particle to particle interactions:

Theorem 4.1 (Multipole Expansion). *Let s be a node in the quad-tree centred at $\mathbf{s} = (s_1, s_2)$. Assume s is not a neighbour of B . Then the potential $\Phi(\mathbf{x})$ due to s for $\mathbf{x} \in B$ is given by the multipole expansion*

$$\Phi(\mathbf{x}) = \sum_{\mathbf{x}_i \in B} q_i K_0 \left(\frac{|\mathbf{x} - \mathbf{x}_i|}{\alpha} \right) = \sum_{l=-\infty}^{\infty} M_l K_l \left(\frac{|\mathbf{x} - \mathbf{s}|}{\alpha} \right) e^{il\theta_{\mathbf{x}}},$$

where

$$M_l = \sum_{\mathbf{x}_i \in B} q_i I_l \left(\frac{|\mathbf{x}_i - \mathbf{s}|}{\alpha} \right) e^{-il\theta_{\mathbf{x}_i}}.$$

Here, I_l and K_l are the l^{th} order modified Bessel's function of the first and second kind, respectively, and $\theta_{\mathbf{x}}$ and $\theta_{\mathbf{x}_i}$ denote the polar angles for \mathbf{x} and \mathbf{x}_i with respect to \mathbf{s} .

The second step of the FMM is to construct a *local expansion* for each node B . In order to translate the multipole expansions to local ones, we apply the following result:

Theorem 4.2 (Local Expansion). *Suppose a multipole expansion associated with node s centred at \mathbf{s} is given by*

$$\Phi(\mathbf{x}) = \sum_{n=-\infty}^{\infty} M_n K_n \left(\frac{|\mathbf{x} - \mathbf{s}|}{\alpha} \right) e^{in\theta_{\mathbf{x}}}.$$

Further assume s is in the “far field” of B . Then for any $\mathbf{x} \in B$, $\Phi(\mathbf{x})$ can be represented as a local expansion

$$\Phi(\mathbf{x}) = \sum_{l=-\infty}^{\infty} L_l I_l \left(\frac{\rho}{\alpha} \right) e^{il\theta},$$

where (ρ, θ) denote the polar coordinates of \mathbf{x} with respect to the centre of B , and the local coefficients are expressed using the “multipole to local” translation operator defined by

$$L_l = \sum_{n=-\infty}^{\infty} M_n K_{l-n} \left(\frac{\rho_0}{\alpha} \right) e^{-i(l-n)\theta_0}.$$

Here, (ρ_0, θ_0) are the polar coordinates of the centre of B with respect to \mathbf{s} .

The work required in the above Theorem represents the bulk of the operation count in the FMM. This work can be significantly reduced by using plane wave representations and a “diagonal translation operator” (c.f. Appendices A and C in [4]). The savings in cost are shown in table 1. The rest of the original FMM procedure proceeds as in sections 2 and 3 in [4], and we will not elaborate further.

N	Direct	Original FMM	Plane-Wave FMM
64	0.01	0.06	0.06
128	0.03	0.08	0.08
256	0.08	0.22	0.12
512	0.30	0.47	0.21
1024	1.15	1.03	0.36
2048	4.59	2.02	0.60
4096	18.26	3.86	1.36
8192	73.42	9.46	2.02
16384	292.52	14.80	4.60
32768	1168.22	38.44	7.04
65536	4689.96	56.55	17.42

Table 1: A comparison of the CPU time (in seconds) required to compute the potential due to a set of points directly, via the original FMM based on multipole expansions, or via the FMM based on plane-wave expansions. Half of the points are randomly placed in the unit square and the other half are concentrated on two elliptical boundaries.

Evaluating (9), however, is not as simple as calculating a sum in the form of (10). In the case of the Neumann problem, we are required to evaluate a potential of the form

$$\begin{aligned} \frac{\partial \Phi}{\partial n_{\mathbf{x}_j}}(\mathbf{x}_j) &= \sum_{i \neq j} q_i \nabla_{\mathbf{x}_j} K_0 \left(\frac{|\mathbf{x}_j - \mathbf{x}_i|}{\alpha} \right) \cdot \mathbf{n}_j \\ &= \nabla_{\mathbf{x}_j} \Phi(\mathbf{x}_j) \cdot \mathbf{n}_j. \end{aligned}$$

This evaluation is relatively straightforward as $\nabla_{\mathbf{x}_j} \Phi(\mathbf{x}_j)$ relies on the same multipole coefficients as $\Phi(\mathbf{x}_j)$. In the case of the Dirichlet problem, we are required to evaluate a potential of the form

$$\frac{\partial \Phi}{\partial n_{\mathbf{x}_i}}(\mathbf{x}_j) = \sum_{i \neq j} q_i \nabla_{\mathbf{x}_i} K_0 \left(\frac{|\mathbf{x}_j - \mathbf{x}_i|}{\alpha} \right) \cdot \mathbf{n}_i.$$

In this case, the multipole coefficients do change:

$$M_l = \sum_{\mathbf{x}_i \in B} q_i \frac{\partial}{\partial n_{\mathbf{x}_i}} I_l \left(\frac{|\mathbf{x}_i - \mathbf{s}|}{\alpha} \right) e^{-il\theta_{\mathbf{x}_i}}.$$




Figure 2: *The solution to Example 1.*

5. Numerical Results

The algorithms described above have been implemented in Fortran. The tolerance for convergence of GMRES is set to 10^{-11} . Here, we illustrate the performance on a variety of examples. All timings cited are for a Mac Pro 2.1 with two 3GHz Quad-Core Intel Xeon processors.

Example 1: We first consider the problem of solving (2) with $\alpha = 0.1$ in a bounded, circular domain with ten interior elliptic contours (see figure 2). We generate the Dirichlet boundary conditions from

$$u(\mathbf{x}) = \sum_{k=1}^{10} K_0 \left(\frac{|\mathbf{x} - \mathbf{x}_k|}{\alpha} \right), \quad (11)$$

where \mathbf{x}_k is a point inside Γ_k . We test the performance of our methods using quadrature rules of varying degree of accuracy. The results are shown in table 2 through table 6. These tables confirm that the number of GMRES iterations required for convergence is independent of N . Also, we see near-linear scaling of the CPU time with N . In terms of overall accuracy, we see the best overall performance with the $O(h^8 \log h)$ scheme.

N	# Iterations	CPU Time	Error
704	45	12.1	$3.80 \cdot 10^{-6}$
1408	45	20.3	$3.36 \cdot 10^{-7}$
2816	45	40.1	$4.20 \cdot 10^{-8}$
5632	45	83.9	$5.25 \cdot 10^{-9}$

Table 2: Performance of the algorithm on example 1. Trapezoid Rule.

N	# Iterations	CPU Time	Error
704	45	12.3	$9.95 \cdot 10^{-7}$
1408	45	20.2	$3.60 \cdot 10^{-7}$
2816	45	41.1	$5.71 \cdot 10^{-8}$
5632	45	87.7	$8.65 \cdot 10^{-9}$

Table 3: Performance of the algorithm on example 1. Quadrature is $O(h^2 \log h)$.

N	# Iterations	CPU Time	Error
704	45	12.5	$1.11 \cdot 10^{-6}$
1408	45	21.1	$4.92 \cdot 10^{-10}$
2816	45	42.3	$3.19 \cdot 10^{-12}$
5632	45	88.9	$2.26 \cdot 10^{-13}$

Table 4: Performance of the algorithm on example 1. Quadrature is $O(h^4 \log h)$.

N	# Iterations	CPU Time	Error
704	45	13.3	$1.11 \cdot 10^{-6}$
1408	45	22.5	$5.41 \cdot 10^{-10}$
2816	45	45.0	$5.88 \cdot 10^{-13}$
5632	45	94.2	$8.55 \cdot 10^{-14}$

Table 5: Performance of the algorithm on example 1. Quadrature is $O(h^8 \log h)$.

N	# Iterations	CPU Time	Error
704	45	14.7	$1.11 \cdot 10^{-6}$
1408	45	25.2	$5.50 \cdot 10^{-10}$
2816	45	50.2	$1.40 \cdot 10^{-11}$
5632	45	104.5	$1.44 \cdot 10^{-11}$

Table 6: Performance of the algorithm on example 1. Quadrature is $O(h^{16} \log h)$.

Example 2: We now consider an exterior Neumann problem, whose boundary conditions are, again, derived from (11). The geometry is also the same as in the preceding example, minus the outer boundary. Based on previous results, we select the $O(h^8 \log h)$ quadrature rule, and we examine the performance of the methods for different values of α in (2). The results are shown in table 7 through table 10. Interestingly enough, the conditioning of the linear system appears relatively independent of α for $\alpha > 0.1$. This is not the case for the Dirichlet problem, where we observed that the number of GMRES iterations visibly increases as α increases in size. This is to be expected, as the governing equation becomes more harmonic, and the integral equation based on the double layer potential for Laplace’s equation in multiply connected domains is known to become rank deficient [5]. Since we are anticipating applications in which α is proportional to a time step, we are not concerned here with this behaviour.

N	# Iterations	CPU Time	Error
704	23	4.3	$1.99 \cdot 10^{-6}$
1408	23	6.8	$2.08 \cdot 10^{-11}$
2816	23	13.5	$1.71 \cdot 10^{-13}$
5632	23	26.7	$1.81 \cdot 10^{-13}$

Table 7: Performance of the algorithm on example 2 with $\alpha = 10$.

Example 3: We next consider a larger-scale problem. We compute the solution to the modified Helmholtz equation with $\alpha = 0.1$ in a bounded domain with 100 elliptical contours that have varying proportions and alignment (figure 3). On each of these contours, we prescribe a constant Dirichlet boundary conditions selected randomly from $(-1, 1)$. Each contour is discretized with 512 points, resulting in a matrix of order 51712. The total CPU time to solve

N	# Iterations	CPU Time	Error
704	23	4.5	$1.83 \cdot 10^{-6}$
1408	23	6.9	$1.87 \cdot 10^{-11}$
2816	23	13.7	$4.37 \cdot 10^{-13}$
5632	23	26.7	$6.33 \cdot 10^{-13}$

Table 8: Performance of the algorithm on example 2 with $\alpha = 1$.

N	# Iterations	CPU Time	Error
704	25	6.5	$5.43 \cdot 10^{-7}$
1408	25	8.8	$6.91 \cdot 10^{-12}$
2816	25	16.4	$4.67 \cdot 10^{-12}$
5632	25	31.6	$4.21 \cdot 10^{-12}$

Table 9: Performance of the algorithm on example 2 with $\alpha = 0.1$.

N	# Iterations	CPU Time	Error
704	15	3.6	$6.33 \cdot 10^{-11}$
1408	15	6.9	$6.91 \cdot 10^{-11}$
2816	15	14.5	$2.77 \cdot 10^{-12}$
5632	15	27.0	$6.45 \cdot 10^{-12}$

Table 10: Performance of the algorithm on example 2 with $\alpha = 0.01$.




Figure 3: *The solution to example 3. There are 512 points per hole, $N = 51712$, 88 GMRES iterations are required, consuming 1510.7 s of CPU time. The plot on the right shows a close-up of the solution from a small region from the upper left quadrant of the domain.*

this linear system to a tolerance of 10^{-11} required approximately 25 minutes.

Example 4: While the domains in the preceding examples are multiply connected, the individual contours have consisted of rather simple shapes. In this example, we consider more complex contours (figure 4). We solve equation (2) in an unbounded domain with $\alpha = 0.1$. Dirichlet boundary conditions are assigned in the following way: we assign a value of 1 on the wheels of the truck, a value of 2 to the rest of the truck and a value of 3 on the person.

6. Conclusions

We have presented a class of integral equation methods for the solution to the modified Helmholtz equation in bounded or unbounded multiply-connected domains. Using a fast-multipole accelerated iterative method, our solution procedure requires $O(N)$ operations, where N is the number of nodes in the discretization of the boundary. With these techniques, and using only modest computational resources, we were able to accurately and efficiently solve large-scale problems in complex domains.

The immediate focus of future work is to develop integral equation techniques to solve (1) in arbitrary geometry. To do this, we must couple our solver with the solver presented in [4]. Once that is achieved, we will apply these integral equation techniques to solve the equations that arise from the temporal discretization of convection-diffusion type equations. In this way,




Figure 4: The solution to example 4 with $N=4740$. 42 GMRES iterations are required, requiring 36.4 s of CPU time.

we hope to demonstrate that integral equation methods offer an attractive alternative to conventional finite element and finite difference methods for many large-scale problems from engineering and physics.

- [1] M. Abramowitz and I.A. Stegun, editors. *Handbook of Mathematical Functions*. Dover, New York, NY, 1964.
- [2] B. Alpert. Hybrid Gauss-Trapezoidal Quadrature Rules. *SIAM Journal of Scientific Computing*, 20:1551–1584, 1999.
- [3] J. Carrier, L. Greengard, and V. Rokhlin. A fast adaptive multipole algorithm for particle simulations. *SIAM J. Sci. Statist. Comput.*, 9:669–686, 1988.
- [4] H. Cheng, J. Huang, and Leiterman T. An Adaptive Fast Solver for the Modified Helmholtz Equation in Two Dimensions. *J. Comp. Phys.*, 211:616–637, 2006.
- [5] L. Greenbaum, A. Greengard and G.B. McFadden. Laplace’s Equation and the Dirichlet-Neumann Map in Multiply Connected Domains. *J. Comp. Phys.*, 105:267–278, 1993.

- [6] L. Greengard and J. Huang. A new version of the fast multipole method for screened Coulomb interactions in three dimensions. *J. Comp. Phys.*, 180:642–658, 2002.
- [7] L. Greengard and M.C. Kropinski. An Integral Equation Approach to the Incompressible Navier-Stokes Equations in Two Dimensions. *SIAM Journal of Scientific Computing*, 20:318–336, 1998.
- [8] L. Greengard, M.C. Kropinski, and A. Mayo. Integral Equation Methods for Stokes Flow and Isotropic Elasticity in the Plane. *J. Comp. Phys.*, 125:403–414, 1996.
- [9] L. Greengard and J.-Y. Lee. A Direct Adaptive Poisson Solver of Arbitrary Accuracy. *J. Comp. Phys.*, 125:415–424, 1995.
- [10] L. Greengard and V. Rokhlin. A new version of the fast multipole method for the Laplace equation in three dimensions. *Acta Numer.*, 6:229–269, 1997.
- [11] L. Greengard and Rokhlin V. A Fast Algorithm for Particle Simulations. *J. Comp. Phys.*, 73:325–348, 1987.
- [12] Leslie Greengard, Denis Gueyffier, Per-Gunnar Martinsson, and Vladimir Rokhlin. Fast direct solvers for integral equations in complex three-dimensional domains. *Acta Numer.*, pages 243–275, 2009.
- [13] S. Kuo, B. Tidor, and J. White. A meshless, spectrally accurate, integral equation solver for molecular surface electrostatics. *ACM JETC*, 4:353–367, 2008.
- [14] B. Lu, X. Cheng, J. Huang, and J.A. McCammon. An adaptive fast multipole boundary element method for poisson–boltzmann electrostatics. *J. Chem. Theory and Comp.*, 5:1692–1699, 2009.
- [15] B.Z. Lu, Y.C. Zhou, J.J. Holst, and J.A. McCammon. Recent progress in numerical methods for the poisson-boltzmann equation in biophysical applications. *Comm. Comp. Phys.*, 3(5):973–1009, 2008.
- [16] P.G. Martinsson and V. Rokhlin. A fast direct solver for boundary integral equations in two dimensions. *J. Comp. Phys.*, 205:1–23, 2005.

- [17] N. Nishimura. Fast multipole accelerated boundary integral equation methods. *Appl. Mech. Rev.*, 55(4):299–324, July 2002.
- [18] Y. Saad and M.H. Schultz. Gmres: a generalized minimum residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7:856–869, 1986.
- [19] Lexing Ying, George Biros, and Denis Zorin. A high-order 3d boundary integral equation solver for elliptic pdes in smooth domains. *J. Comp. Phys.*, 219:247–275, 2006.